

08/12/04

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application No.: 10/692,667

APPLICANT: CHANDRASEKHARA NILAKANTAN, ET AL.

Filed: 10/24/2003

TC/A.U.: 2183

Examiner: Not Yet Assigned

Docket No.: 6728P001

Customer No.: 08791

Confirmation No.: 9960

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail with sufficient postage in an envelope addressed to the Commissioner for Patents, PO Box 1450, Alexandria, VA 22313-1450

on August 10, 2004
Date of Deposit

Jennifer L. Stewart
Name of Person Mailing Correspondence

August 10, 2004
Signature Date

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

PETITION TO MAKE SPECIAL (37 CFR 1.102(d))

Sir or Madam:

Applicants hereby petition to make this new application special. This application has not received any examination on the merits.

(A) FEE

Applicants hereby enclose a check in the amount of \$130.00 for the petition fee required by 37 C.F.R. § 1.17(h). Furthermore, the Commissioner is hereby authorized to charge payment of any fee due under 37 C.F.R. § 1.16 and § 1.17 associated with this communication or any future communication in this or any related application filed pursuant to 37 C.F.R. § 1.53 or credit any overpayment to Deposit Account No. 02-2666.

(B) CLAIMS

Either (1) all pending claims in this application are directed to a single invention, or (2) if the Office determines that all the claims are not obviously directed to a single invention, applicants will make an election without traverse in response to notification under the established telephone restriction practice.

08/12/2004 JBALINAN 00000133 10692667

01 FC:1460

130.00 OP

(C) SEARCH

A search for relevant prior art was made and the fields of search included:

U.S. patents and published applications in classes/subclasses:

370/469, 370/465, 370/463, 70/428, 370/423, 370422, 370/411, 370/401
370/400, 370/235, 370/232, 370/231, 370/230, 370/219, 370/217, 709/250,
709/249, 709/241, 709/239, 709/238, 370/237, 709/232, 709/229, 709/227,
709/226, 709/203, 709/105,

publications; and

foreign patents and published applications.

(D) COPIES OF REFERENCES / INFORMATION DISCLOSURE STATEMENT

Attached are copies of references located during the above-referenced search that are deemed most closely related to the subject matter encompassed by the claims. Each of these references is listed in the attached Information Disclosure Statement. Applicants respectfully request that all references be considered and entered into the record of the present application.

The submission of these references is for the purpose of providing a complete record and is not a concession that the references listed therein are prior art to the invention claimed in the patent application. The right is expressly reserved to establish an invention date earlier than the above-identified filing date in order to remove any reference submitted herewith as prior art should it be deemed appropriate to do so.

Further, the submission of the references is not to be taken as a concession that any reference represents art that is relevant or analogous to the claimed invention. Accordingly, the right to argue that any reference is not properly within the scope of prior art relevant to an examination of the claims in the above-identified application is also expressly reserved.

This Information Disclosure Statement is being filed before the mailing date of a first Office Action on the merits. Therefore, Applicants believe no fee is due; however, should a fee be due, the Commissioner is hereby authorized to charge Deposit Account No. 02-2666.

(E) DETAILED DISCUSSION OF THE REFERENCES

A detailed discussion of the references deemed most closely related to the subject matter encompassed by the claims is provided below.

Each selected reference fails to anticipate the present invention as claimed. To anticipate

a claim, the reference must teach every element of the claim. "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987).

Furthermore, the selected references fail to establish a prima facie case of obviousness because the references, individually or in combination, neither teach nor suggest all the claim elements and limitations required by the patent application. Moreover, there is no motivation or suggestion in these references for their combination; and even assuming there were such motivation or suggestion, no combination of these references teaches or suggests the invention as claimed.

To establish a prima facie case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991).

Therefore, it is submitted that all pending claims are distinguishable over the cited references, taken alone or in combination, and should be allowed.

Claims 1, 29, 46, and 59 of the present application

The present invention relates to systems and methods for enabling and automating the delivery of interactive television content to subscribers. The following are the independent claims of the present application:

1. A method, comprising determining, in a centralized fashion, paths for flows within a multi-stage network made up of ~~clusters of processing nodes~~ grouped into clusters having similar functionality, and encoding node selection information representing flow path decisions for all clusters of the multi-stage network in packets of each flow within the multi-stage network.
(Amended)

15. A method, comprising determining, in a distributed fashion, paths for flows within a multi-stage network made up of ~~clusters of processing nodes~~ grouped into clusters having similar functionality, and encoding node selection information representing flow path decisions for all clusters of the multi-stage network in packets of each flow within the multi-stage network.

(Amended)

23. A method, comprising replicating, at a node of a multi-stage network in which nodes are grouped into clusters of nodes having similar application functionality, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced.

26. A system comprising a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern.

49. A method, comprising establishing a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices.

Wong et al. (US 2004/0037278), (US 26614756), (US 6363077)

The invention of Wong'278, Wong'756 and Wong'077 are directed to a "communication network switch that provides a method and apparatus for balancing the loading of aggregated network links of the trunk, thereby increasing the data transmission rate through the trunk" (Abstract). The method uses a packet routing unit and a load balancing unit to determine the data path. Specifically, Wong'278, Wong'756 and Wong'077 describe that when a packet routing unit receives a packet, it determines a destination port value indicating the trunked destination port, based on the source MAC address and destination MAC address of the packet. The load

balancing unit then elects on the destination port from the various ports of the trunk destination port (Wong'278, Page 3, [0039]; Wong'756, Col.5, lines 29-51; Wong'077, Col.5, lines 29-51).

The data path selection method as taught by Wong'278, Wong'756 and Wong'077 clearly fails to teach or suggest each and every element of the present claims:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing *a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network*, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that

emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Applicants respectfully submit that in light of the above arguments, Wong' 278, Wong'756 and Wong'077 have been overcome.

Beshai (US 2003/0189947)

Beshai'947 provides a method that defines "a protocol that uses an adaptive packet header to simplify packet routing and increase transfer speed" (Abstract). However, the routing method as taught by Beshai'947 is based on "a function of both the static cost and route vacancy" (Page 10, [0130]). Indeed, the path selection is significantly different from the present claims:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

In addition, Beshai'947 does not teach or suggest the "virtual connectivity grid" of claim 26 and 49 of the present invention. Accordingly, the present claims are patentable over Beshai.

Sarkinen et al. (US 2003/0058880)

Sarkinen'880 presents a switching method that includes “providing virtual input and output queues coupled to a switch fabric comprising switch elements; collecting statistics regarding the virtual input and output queues; and controlling packet queuing for the virtual input and output queues using the collected statistic to provide congestion control for the virtual input and output queues and the switch fabric” (Claim 1).

Though Sarkinen'880 teaches virtual input and output queues, he fails to anticipate the present invention wherein:

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Indeed, Sarkinen'880 is only concerned with the virtual input and output queues of a switch. In contrast, the present invention teaches a “virtual connectivity grid” of a network system. This alone is sufficient for the present claims to be patentable over Sarkinen'880.

In addition, Sarkinen'880 fails to teach or suggest that the processing nodes are “grouped into clusters having similar functionality” (Claim 1, 15, 23). Therefore, the present claims are patentable over Sarkinen'880.

Shinomiya et al. (US 2003/0037165)

Shinomiya'165 is directed towards a dynamic load sharing system using a virtual router. Shinomiya'165 discloses that the system includes "a plurality of equipment units each functioning as a router which constitutes a virtual router having a single common address; and end systems being connected to the network through the virtual router. Among the plurality of routers constituting the virtual router, one equipment unit functioning as a router is assigned as a master router, while the other equipment unit(s) is assigned as a backup router. The master router dynamically allocates packet condition for defining the routing object to each router, then to advertise to the backup router. Routing processing between the network and the end system is performed by the plurality of routers each having a routing function" (Abstract).

Clearly, Shinomiya'165 is only concerned with the dynamic load sharing system of the routers. In particular, Shinomiya'165 fails to disclose the present claims which include (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid (Claims 26 & 49). Consequently, the present claims are patentable over Shinomiya'165.

Basturk (US 2003/0023750)

Basturk'750 provides a method for controlling data flow over a data-packet-network according to specific destinations. In addition, Basturk'750 discusses a method for altering the established course of a data path. The method includes the steps of " (a) accessing a particular node in the course of the data path toward the destination; (b) delivering to the node at least one cost parameter specific to the destination to replace at least one existing cost parameter specific to the destination reported by the node; (c) installing the at least one new parameter replacing the at least one old parameter; and (d) reporting the at least one new parameter from the affected node to nodes neighboring the affected node such that data traffic routed to the specific destination assumes an altered route to the destination" (Basturk'750, Claim 1).

Indisputably, Basturk'750 is only concerned with using cost values to determine the data path. He fails to teach or suggest the present claims. In particular, he does not disclose (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid (Claims 26 & 49). Therefore, the present claims are patentable over Basturk'750.

Durinovic-Johri et al. (US2002/0176359)

Durinovic-Johri'359 describes an apparatus for managing packet flow in network routers. In particular, the apparatus enables routers to "communicate the congestion status among the ports inside routers in the network, and substantially eliminates packet dropping due to congestion by providing overflow paths for destination IP addresses. Each router in a network stores at least two possible output paths for selected destination IP addresses, so that the router may direct the output of packets appropriately when congestion is detected on one of the paths. A forwarding table stores the possible output paths for each destination IP address (Abstract)".

At the protocol level, Durinovic-Johri'359 teaches using a forwarding or routing table. In particular, "at least two possible paths are selected by the routing protocol for each destination IP address and their first hops or outgoing links are stored in the forwarding table of the router" (Page 2, [0023]).

This is significantly different from the present claims. Indeed, Durinovic-Johri'359 fails to teach or suggest any of the following:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks*

established within the virtual connectivity grid, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Accordingly, the present claims are patentable over Durinovic-Johri'359.

Ayyagari et al. (US2002/0101822)

Ayyagari'822 presents a new protocol for making scheduling, routing and access control decisions by integrating the characteristics of the physical layer, the link layer and the network layer in a networked computer environment that maximizes QoS. (Page 2, [0022]). Ayyagari'822 further discloses that the network nodes "are organized into at least one of a cluster and a clique and the network has a network-wide capacity to send data packets in slots delineating time frames on the network between the nodes on the network defined by the links" (Page 2, [0023]). However, Ayyagari'822 fails to further define how are nodes are grouped together as a cluster or a clique. Indeed, Ayyagari'822 only discloses that "it will be understood that clusters 68 and cliques 66 are well-known in ad-hoc wireless networks 60 and the nodes 62 thereof are organized into clusters/cliques by a known algorithm" (Page 5, [0061]).

This is significantly different from the present claims where the nodes are grouped into clusters having "similar application functionality"(Claims 1, 15 & 23). In addition, Ayyagari'822 does not disclose any of the following:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of processing nodes grouped into clusters having similar functionality, and encoding node selection information representing flow path decisions for all clusters of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a *virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network*, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Ayyagari'822 certainly fails to teach each and every element of the present claims. Therefore, the present claims are not anticipated by Ayyagari'822.

Bak et al. (US2002/0083194)

The invention of Bak'194 relates to a routing method for traffic load distribution in a packet-switched network. The method includes "a first step for computing an average cost

between the start node and a plurality of nodes stored in a routing table and selecting a node having a cost lower than an average cost value K as a candidate of an intermediate node; a second step for randomly selecting one among the intermediate node candidates when a packet to be transmitted from the start node is generated and determining the selected node as an intermediate node of the packet; a third step for storing information with respect to a path setting bit(b) a destination node address and an intermediate node address in a header region of the is packet and transmitting the packet to the intermediate node through the path of the lowest cost using an information stored in the routing table; and a fourth step for transmitting the packet to the destination node through the lowest cost path using the information stored in the routing table when the packet arrives at the intermediate node” (Bak’194, Claim 1).

However, Bak’194 fails to teach or suggest each and every element of the present invention. Specifically, Bak’194 does not disclose any of the following:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being

established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Indisputably, Bak'194 fails to teach each and every element of the present claims, and, as such, the present invention is patentable over Bak'194.

Swildens (US2001/0034792)

Swildens'792 provides a Decrypting Load Balancing Array system that "uses a Pentaflow approach to network traffic management that extends across an array of Decrypting Load Balancing Array (DLBA) servers sitting in front of back end Web servers. One of the DLBA servers acts as a scheduler for the array through which all incoming requests are routed. The scheduler routes and load balances the traffic to the other DLBA servers (including itself) in the array. Each DLBA server routes and load balances the incoming request packets to the appropriate back end Web servers. Responses to the requests from the back end Web servers are sent back to the DLBA server which forwards the response directly to the requesting client. SSL packets are decrypted in the DLBA server before being routed to a back end Web server, allowing the DLBA server to schedule SSL sessions to back end Web servers based on a cookie or session ID. Response packets are encrypted by the DLBA server before being forwarded to the client" (Abstract).

The routing approach of Swildens'792 clearly uses only the DLBA server. This teaches away from the present invention wherein the routing method includes (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid (Claims 26 & 49). Indisputably, the present claims are patentable over Swildens'792.

Kanekar et al. (US6751191)

Kanekar'191 relates to load sharing and redundancy in a network. The approach uses "a master router and a slave router operating in the same chassis and having a shared set of interfaces. Prior to failure of the master router, the master router communicates shared state information to the slave router. In addition, the slave router operates in 'standby mode' to obtain information from the shared set of interfaces" (Abstract). This is achieved by "assigning a shared IP address and a shared MAC address to both a first router and a second router so that the shared IP and MAC addresses are shared between the first router and the second router. Additionally, a first MAC address is assigned to the first router and a second MAC address is assigned to the second router. The default gateway is configured on the hosts such that a default gateway IP address is associated with the shared IP address. The shared IP and MAC addresses are associated with one of the routers (e.g., the first router or master router). When the master fails, the slave takes over both the shared IP address and the shared MAC address" (Abstract).

As presented above, Kanekar'191 designs a master router and a slave router to serve as a backup for each other through a shared IP address. This approach is nowhere close to the present invention which includes (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid (Claims 26 & 49). As such, the present claims are patentable over Kanekar'191.

Aviani et al. (US6742044)

Aviani'044 presents a distributed network traffic load balancing technique without using a gateway router. The technique presented by Aviani includes an intercept server that "sits in front of a host server, and intercepts packets routed to the host server. When desired, packets which are intercepted by the intercept server are replicated, encapsulated and tunneled to selected client servers in the overlay network. The tunneled packets are received and processed by each of the selected client servers, whereupon each of the selected client servers generates a respective spoofed response to the source device identified in the header of the originally intercepted packet. Further, according to the technique of the present invention, each of the selected client servers transmits its respective spoofed response to the identified source device at substantially the same time. The client server associated with the spoofed response which is first received at the identified source device is considered to have the relatively shortest propagation delay to the identified source device, and is identified as the successful client server. Thereafter, the source

device will be directed or redirected to communicate directly with the successful client server when subsequently attempting to access information from the host server” (Abstract).

The data path selection of Aviani’044 is clearly determined by the response time from the client server to the source device. The client response with the relatively shortest propagation delay indicates the shortest path and, therefore, is the preferred data route. This is significantly different from the present invention wherein:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing *a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network*, each of the computer systems being communicatively coupled to respective ports of one or more

physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Aviani'044 fails to teach each and every element of the present claims and, as such, the present claims are patentable over Aviani'044.

Devarakonda et al. (US6424992)

Devarakonda'992 teaches an affinity-based router and routing method. In his approach, a TCP router is designed to maintain affinity tables and records to indicate which node a client was routed to. The affinity is due to the state at the server either due to previous routing request or data affinity at the server. (Abstract).

Devarakonda'992 is only concerned with using affinity information to determine a data path. This is fundamentally different from the present invention which includes (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid for connecting network systems (Claims 26 & 49). Applicants respectfully submit that in light of the above arguments, Devarakonda'992 has been overcome.

Lin et al. (US6272522)

The invention of Lin'522 presents a data packet switching and server load balancing device using a general-purpose multiprocessor computer system. The data routing of Lin'522 includes switching processors that "route received ones of the data packets to a selected one of the external networks in accordance with information included in a header portion of the data packets and the load distribution configuration data" (Abstract). The load distribution configuration data is generated through "complex calculations on the raw system load information" (Col.6, lines 1-3).

The routing technique as taught by Lin'522 clearly fails to anticipate the present invention which includes (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid for connecting network systems (Claims 26 & 49).

As Lin'522 fails to teach or suggest each and every element of the present claims, the present claims are patentable over Lin'522.

Basilico (US6243360)

Basilico'360 presents a network server system having dynamic load balancing of messages in both inbound and outbound directions. In particular, Basilico'360 discloses that when workstations sends a data packet including network destination addresses to a network switch, a "header is prepended to the data packet, the header identifying a switch output or destination port for transmitting the data packet to the network destination address" (Abstract). The header is generated from "the destination address via the CAM/RAM lookup or other switch mapping method" (Col.6, lines 1-3).

Though Basilico'360 teaches prepending a header to a data packet, he still fails to teach or suggest each and every element of the present invention. Indeed, Basilico'360 does not disclose the following:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a virtual connectivity grid overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces*

so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

The routing approach of Basilico'360 is significantly different from the present claims. Consequently, Basilico'360 fails to anticipate the present claims.

Walker (US5613069)

Walker'069 relates to a non-blocking packet switching network with dynamic routing codes having incoming packets diverted and temporarily stored in processor input when network output is not available. As illustrated in Figure 2, "if user wants to establish a call to user 3b, user 3a initially sends a packet requesting routing information without routing information. Switch 2a detects the absence of routing information and passes the packet to name management switch 5a, which identifies the destination address, held in the packet payload, and having knowledge of the network topology establishes a route through the network from user 3b to user 3b. This information is incorporated in the packet payload, which is returned to user 3a using information in the packet trailer, as will be described, to determine the return path. When user 3a receives the returned packet containing the routing information, this can be incorporated in the header of subsequent packets to send them to their destination." (Col.6, line 62 – Col.7, line 8)

Though Walker'069 discloses incorporating routing information in the header of data packets, he still fails to teach or suggest each and every element of the present claims. In particular, the following features of the present claims distinguish the present invention from Walker'069:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

Indeed, Walker'069 is absolutely unconcerned with organizing network nodes into cluster of similar application functionality for data path determination.

In addition, Walker'069 specifically states that “the use of routing codes, rather than a link id. (virtual channel, virtual path), as means of routing packets means that the routing codes can be interpreted directly by hardware in each packet switch”. This clearly teaches away from the present claims, which read:

a *virtual connectivity grid* overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a *virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network*, each of the computer systems being communicatively coupled to respective ports of one or more

physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Walker'069 clearly prefers "hardwired switching logic" (Col.5, lines 3-7). Indisputably, this contradicts with the present claims wherein "virtual connectivity grid" is used.

Walker'069 fails to teach or suggest each and every element of the present claims. Furthermore, Walker'069 teaches away from the present claims. Therefore, the present claims are patentable over Walker'069.

Waclawsky et al. (US5495426)

Waclawsky'426 discloses a dynamic realtime, inband routing mechanism that uses an Event Driven Interface (EDI). The EDI is programmed with control vectors to identify load balancing and load distribution bit patterns in the data frames on the data communications network. In his approach, each data packet has a flag field with path characteristic pattern and bit capacity to match the maximum number of potential paths that a packet can take (Col.2, lines 32-44). Because "the EDI is pre-configured with all possible path alternatives, only a specific path characteristic pattern is required to enable the EDI to get the packet through the communications network" (Col.2, lines 34-37).

As established above, the routing method as suggested by Waclawsky'426 features path characteristic pattern. This is absolutely different from the present claims. The present claims include the following features that are not taught or suggested by Waclawsky'426:

determining, in a centralized fashion, paths for flows within a multi-stage network made up of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* of the multi-stage network in packets of each flow within the multi-stage network. (Claim 1)

determining, in a distributed fashion, paths for flows within a multi-stage network made up of clusters of *processing nodes grouped into clusters having similar functionality*, and *encoding node selection information representing flow path decisions for all clusters* in packets of each flow within the multi-stage network. (Claim 15)

replicating, at a node of a multi-stage network in which *nodes are grouped into clusters of nodes having similar application functionality*, an entire flow, in both directions, and designating a flow path for a resulting replicated flow that is different than an original flow path for an original flow from which the replicated flow was produced. (Claim 23)

a *virtual connectivity grid* overlaid on a physical network in which nodes are coupled to one or more physical switches through respective physical interfaces, *the virtual connectivity grid including virtual interfaces overlaid over the node physical interfaces so as to provide communicative coupling of the nodes to one or more virtual networks established within the virtual connectivity grid*, the communicative coupling being established by virtual links arranged so as to emulate physical connections in a desired connectivity pattern. (Claim 26)

establishing a *virtual connectivity grid configured to permit arbitrary interconnections among a first number of computer systems within a computer network*, each of the computer systems being communicatively coupled to respective ports of one or more physical network switching devices, through a second number of virtual links that emulate physical network connectivity mechanisms as a result of configurations of one or more virtual networks (VLANs) overlaid on ports of the physical network switching devices. (Claim 49)

Indeed, Waclawsky'426 is indifferent about (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar application functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid for connecting network systems (Claims 26 & 49). Consequently, all of the present claims are patentable over Waclawsky'426.

Raahemi et al., "A load-balancing adaptive routing algorithm in k-ary n-cube interconnection network", CCECE, Vol.2, Page 725-8, Montreal, Canada, 2003.

In this paper, Raahemi presents an Order Preserving Routing algorithm (OPRA) which is an adaptive algorithm based on dimension-order routing and utilizes path diversity available in the network. The OPRA routing algorithm "randomly selects a minimal path from a set of

minimal routes available between source and destination” (Abstract). In addition, the “path randomization is carried out in such a way that the order of packets belonging to the same flow is not changed, i.e. always the same path is selected for specific flow” (Page 727, Paragraph 2)

However, Raahemi fails to anticipate the present invention which includes (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar application functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid for connecting network systems (Claims 26 & 49). Raahemi does not teach or suggest each and every element of the present claims, and accordingly, all the present claims are patentable over Raahemi.

Abrahamsson et al., “A multi-path routing algorithm for IP networks based on flow optimization”, QofIS/ICQT, Vol.2511, Pages 135-144, Switzerland, 2002.

Abrahamsson relates to an intra-domain routing algorithm based on multi-commodity flow optimization that enables load sensitive forwarding over multiple paths. The approach includes “aggregating the traffic flows destined for the same egress into one commodity in the optimization and using a hash based forwarding mechanism” (Abstract). This results in “computational requirements an order of magnitude smaller than in the traditional models where the problem is modeled with one commodity for each flow from one ingress to one egress node” (Page 144, Paragraph 2)

However, Abrahamsson fails to anticipate the present claims. In particular, Abrahamsson discloses aggregating traffic flow destined for the same egress into one commodity. This is indisputably different from the present claims wherein nodes are “grouped into clusters having similar application functionality” (Claims 1, 15, & 25). Indeed, Abrahamsson is only concerned with the destination address and not the application functionality of the nodes. This alone is sufficient for the claims to be patentable over Abrahamsson.

In addition, Abrahamsson does not teach or suggest using (2) a virtual connectivity grid for connecting network systems (Claims 26 & 49). Unquestionably, Abrahamsson does not teach or suggest each and every element of the present claims, and accordingly, all the present claims are patentable over Abrahamsson.

Chao et al. "A framework for dynamic routing and feedback regulation of packet-switched networks", Telecommunications Network Design and Management, Page 273-301, 2003

Chao presents an approach to modeling, analysis and design for control of traffic in packet-switched networks. In his approach, Chao combines "deterministic network theory, in which link capacities are the dominant network limitation, with queueing theory, in which service rates at nodes are the dominant network limitation." (Page 298, Paragraph 3). The approach relies heavily on analytical techniques from two areas of mathematics – Algebraic topology and Jackson network models in queueing theory.

Nowhere in Chao is there a suggestion or teaching of the present invention which includes (1) determining data path for a multi-stage network made up of processing nodes grouped into clusters having similar application functionality (Claims 1, 15 & 23) and using (2) a virtual connectivity grid for connecting network systems (Claims 26 & 49). As such, all of the present claims are patentable over Chao.

(F) SUMMARY

For at least the foregoing reasons, the claims are patentable over the references located during the above-referenced search that are deemed most closely related to the subject matter encompassed by the claims.

If there are any additional fees associated with this communication, please charge our deposit account 02-2666.

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP



Date: Agust 10, 2004

Chze Koon Chua
Reg. No. 53,831

12400 Wilshire Blvd.
Seventh Floor
Los Angeles, CA 90025
(408) 947-8200